

Network iso-elasticity and weighted α -fairness

Sem Borst, Neil Walton and Bert Zwart

Bell Laboratories, Alcatel-Lucent, 600 Mountain Avenue, Murray Hill, NJ 07974, USA
e-mail: sem@alcatel-lucent.com

University of Amsterdam, Science Park 904, 1098 XH Amsterdam
e-mail: n.s.walton@uva.nl

CWI, Science Park 123, 1098 XG Amsterdam
e-mail: Bert.Zwart@cwi.nl

Abstract: When a communication network's capacity increases, it is natural to want the bandwidth allocated to increase to exploit this capacity. But, if the same relative capacity increase occurs at each network resource, it is also natural to want each user to see the same relative benefit, so the bandwidth allocated to each route should remain proportional. We will be interested in bandwidth allocations which scale in this *iso-elastic* manner and, also, maximize a utility function.

Utility optimizing bandwidth allocations have been frequently studied, and a popular choice of utility function are the weighted α -fair utility functions introduced by Mo and Walrand [17]. Because weighted α -fair utility functions possess this iso-elastic property, they are frequently used to form fluid models of bandwidth sharing networks. In this paper, we present results that show, in many settings, the only utility functions which are iso-elastic are weighted α -fair utility functions.

Thus, if bandwidth is allocated according to a network utility function which scales with relative network changes then that utility function must be a weighted α -fair utility function, and hence, a control protocol that is robust to the future relative changes in network capacity and usage ought to allocate bandwidth in order to maximize a weighted α -fair utility function.

1. Introduction

It is often claimed that both Internet traffic and capacity undergo multiplicative year-on-year increase. Such a rate may not be sustainable over time, even so, it is a pervasive point that the Internet's demand and service rate changes multiplicatively. Given this, in what way should a communication networks resources be allocated in order to scale with such multiplicative effects?

Mo and Walrand [17] introduced a parametrized family of utility functions called the *weighted α -fair utility functions*. A utility function orders the preferences of different network states: A network state is better than a second network state if it has higher utility. When maximized over a network's capacity set, a utility function provides a solution which can be used to allocate network resources. Congestion control protocols such as TCP are argued to allocate bandwidth in this way [15], and congestion protocol models have been designed that provably reach such an operating point [12].

Suppose that a performance improvement is made whereby the network capacity is doubled. Although capacity increases, the criteria by which capacity is evaluated and shared should remain the proportionate. We note that, in particular, the weighted α -fair utility functions obey this scaling property. But are there more? Such a utility function should, also, scale well when the traffic increases. That is, multiplicatively increasing the number of network flows should not alter the preferred allocation of network resources. This is a desirable property: regardless of relative changes, the proportion of network resource each route receives remains the same. Given that the Internet traffic is increasing multiplicatively, it is desirable that each flow has a share of resource that remains proportionate. Once again, the weighted α -fair utility functions will scale in this way. So are the weighted α -fair utility functions the only utility functions that satisfy these scalability properties? We mathematically formulate scaling properties and general settings where this is provably true. We thus provide a theoretical basis to the claim that, if bandwidth is allocated according to a network utility function which scales with relative network changes, then that utility function must be a weighted α -fair utility function, and hence, a control protocol that is robust to the future relative changes in network capacity and usage must to allocate bandwidth in order to maximize a weighted α -fair utility function.

Economists Arrow [1] and Pratt [19] formulated the utility functions that satisfy a certain *iso-elastic property*. In words, an individual is iso-elastic if his preferences enter a bet are unaltered by a multiplicative changes in the bet's stakes and rewards. Arrow and Pratt parametrize the set of iso-elastic utility functions on \mathbb{R}_+ . It is an immediate check that the iso-elastic utility functions are precisely the summands of a weighted α -fair utility function. Thus the reason a weighted α -fair utility function has good scaling properties is because of the iso-elasticity property inherent in its summands. But conversely, if a utility function allocates network resources proportionately then must it be weighted α -fair? The main results of this paper provide settings where scaling properties are equivalent to a utility function being weight α -fair.

The work of Mo and Walrand has received a great deal of attention. This is chiefly because a number of known network fairness criteria correspond to maximizing a weighted α -fair utility function: known Internet equilibria, *TCP fairness* [15], for $\alpha = 2$; the work Kelly [11], weighted proportional fairness, for $\alpha = 1$; the work of Bertsekas and Gallager [2], max-min fairness, for $w = 1$ with $\alpha \rightarrow \infty$ and maximum throughput for $w = 1$ with $\alpha \rightarrow 0$. But, perhaps, the explicit reason for the attention on weighted α -fairness has not been the above iso-elastic property. The iso-elastic property is frequently exploited in order to studying the limit and stability behaviour of associated stochastic network models [3, 9, 16, 10, 7, 21]; some authors have begun to explicitly cite the iso-elastic property [8, 13], and other authors have attempted to derive weighted α -fairness and other fairness criteria from axiomatic properties [14, 20]. Even so, there appears to be little discussion of this iso-elastic scaling property, its equivalence with weighted α -fairness, and what this implies for the performance achieved by a communication network and associated stochastic models. Addressing this

point is the principle contribution of this paper.

The results and sections of this paper are organized as follows. In Section 2, we define the topology and capacity constraints of a communication network. In Section 3, we define what we mean by a network utility function, network utility maximization and we define the weighted α -fair utility functions. In Section 4, we formally define the network iso-elasticity property and the flow scalability property. In Section 5, we prove that a network utility function is network iso-elastic iff it is weighted α -fair. In Section 6, we consider a network topology where each link has some dedicate local traffic flow. On this topology, we show that a network utility maximizing allocation satisfies flow scalability iff it is a weighted α -fair maximizer. In Section 7, in addition to the flow scalability property, we define an access scalability property. We prove the only utility functions that are flow scalable and access scalable are the weighted α -fair utility functions. In proving these results we lay claim to the statement that if bandwidth is allocated according to a network utility function which scales with relative network changes then that utility function must be a weighted α -fair utility function. Thus a control protocol that is robust to the future relative changes in network capacity and usage must allocate bandwidth to maximize a weighted α -fair utility function.

2. Network Structure

We define the topology and capacity constraints of a network. We suppose that there are a set of links indexed by \mathcal{J} . The positive vector $C = (C_j : j \in \mathcal{J})$ gives the capacity of each link in the network. A route consists of a set of links. We let \mathcal{R} denote the set of routes. The topology of the network is defined through a $|\mathcal{J}| \times |\mathcal{R}|$ matrix A . We let $A_{jr} = 1$ if route r uses link j and we let $A_{jr} = 0$ otherwise. We let the positive vector $\Lambda = (\Lambda_r : r \in \mathcal{R})$ denote the bandwidth allocated to each route. At link j these must satisfy the capacity constraint

$$\sum_{j \in \mathcal{J}} A_{jr} \Lambda_r \leq C_j,$$

and thus the set feasible bandwidth allocations is given by

$$\mathcal{C} = \{\Lambda \in \mathbb{R}_+^{\mathcal{R}} : A\Lambda \leq C\}.$$

3. Network Utility

In communication network, we could allocate resources to achieve the highest aggregate throughput. Although a high data rate is achieved, some networks flows maybe be starved. Over the last decade, there has been interest in allocating network resources in a *fair* way. Essentially, fairness is achieved by allocating a positive share of the networks resources to each flow. To achieve this it was recommended to allocate resources inorder to maximize utility, see Kelly [11]. In such a network, each flow expresses its demand via a strictly increasing concave

utility function. A utility function will have a higher slope for smaller throughputs. When maximized, this function can be used express a higher demand for the networks resources. Thus, from this, a form of fairness is achieved.

The average utility of a network flow is

$$U(y, \Lambda) = \frac{1}{\sum_{i \in \mathcal{R}} y_i} \sum_{r \in \mathcal{R}} y_r U_r\left(\frac{\Lambda_r}{y_r}\right).^1 \quad (1)$$

Here \mathcal{R} indexes the routes of a network; y_r gives the number of flows present on route r ; Λ_r gives the flow rate allocated to route r , which is then shared equally amongst the flows present on the route, and finally, U_r gives the utility function of each route r user. We call a utility function of the form (1), a *network utility function*.

We assume throughout this paper that each utility function U_r is increasing, once differentiable and strictly concave.

To the average user, a network utility function orders preferences of different network states, i.e., (y, Λ) is as least as good as $(\tilde{y}, \tilde{\Lambda})$ if $U(y, \Lambda) \geq U(\tilde{y}, \tilde{\Lambda})$. Given the flows present in a network, the best network state is one that maximizes utility. That is, $\Lambda(y)$ the optimum of the *network utility maximization*.

$$\max U(y, \Lambda) \quad \text{over } \Lambda \in \mathcal{C}. \quad (2)$$

Here $\mathcal{C} \subset \mathbb{R}_+^{\mathcal{R}}$ is the set of feasible bandwidth allocations. In general, set of bandwidth allocations will depend on the topology and capacity constraints of a communication network.

The *weighted α -fair* utility functions introduced by Mo and Walrand [17] are given by

$$W_{\alpha, w}(y, \Lambda) := \begin{cases} \frac{1}{\sum_{i \in \mathcal{R}} y_i} \sum_{r \in \mathcal{R}} \frac{w_r y_r}{1-\alpha} \left(\frac{\Lambda_r}{y_r}\right)^{1-\alpha} & \text{if } \alpha \neq 1, \\ \frac{1}{\sum_{i \in \mathcal{R}} y_i} \sum_{r \in \mathcal{R}} w_r y_r \log\left(\frac{\Lambda_r}{y_r}\right) & \text{if } \alpha = 1, \end{cases} \quad (3)$$

for $\Lambda = (\Lambda_r : r \in \mathcal{R}), y = (y_r : r \in \mathcal{R}) \in \mathbb{R}_+^{\mathcal{R}}$ and parametrized by $\alpha > 0$ and $w \in \mathbb{R}_+^{\mathcal{R}}$. As mentioned, the weighted α -fair class has proved popular as it provides a spectrum of fairness criteria which contains proportional fairness ($\alpha = w_i = 1$), TCP fairness ($\alpha = 2, w_i = \frac{1}{T_i^2}$), and converges to a maximum throughput solution ($\alpha \rightarrow 0, w_i = 1$) and max-min fairness ($\alpha \rightarrow \infty, w_i = 1$).

4. Network Iso-elasticity and flow scalability

We now define the two main scaling properties which we will use: network iso-elasticity and flow scalability. Both encapsulate a simple idea, if we multiplicatively change the available network capacity then we will allocate resources proportionately.

¹In this paper, we chose to express the utility of the average network user rather than the aggregate utility of the network. We note the optimum of these functions will be unaffected when maximizing over Λ .

A network utility function ranks the set of network states through the ordering induced by $U(y, \Lambda) \geq U(\tilde{y}, \tilde{\Lambda})$. We say a network utility function is network iso-elastic if this ordering is unchanged by a multiplicative increase in the available bandwidth. More formally,

Definition 1. We define a utility function to be network iso-elastic if, $\forall a > 0$,

$$U(y, \Lambda) \geq U(\tilde{y}, \tilde{\Lambda}) \quad \text{iff} \quad U(y, a\Lambda) \geq U(\tilde{y}, a\tilde{\Lambda}), \quad (4)$$

for each $\Lambda, \tilde{\Lambda}, y, \tilde{y} \in \mathbb{R}_+^{\mathcal{R}}$.

We note that this is equivalent to the expression, where we multiplicatively scale the number of flows, $\forall a > 0$,

$$U(y, \Lambda) \leq U(\tilde{y}, \tilde{\Lambda}) \quad \text{iff} \quad U(ay, \Lambda) \leq U(a\tilde{y}, \tilde{\Lambda}), \quad (5)$$

for each $\Lambda, \tilde{\Lambda}, y, \tilde{y} \in \mathbb{R}_+^{\mathcal{R}}$.

Network iso-elasticity requires that the utility of each network state (y, Λ) scales. What if we only wish the optimal allocation to scale?

Definition 2. We say a utility function optimized on capacity set \mathcal{C} is flow scalable if $\Lambda(y)$ the solution to the optimization problem (2) is such that

$$\Lambda(y) = \Lambda(ay), \quad \forall a > 0. \quad (6)$$

That is the bandwidth allocated to each route is unchanged by multiplicative changes in the number of flows on each route. Making explicit that $\Lambda(y; \mathcal{C})$ is the solution to (2) when optimizing over \mathcal{C} and defining $a\mathcal{C} = \{a\Lambda : \Lambda \in \mathcal{C}\}$. We note that network scalability condition is equivalent to the condition that allocated bandwidth scales proportionately with capacity increases, i.e.,

$$a\Lambda(y; \mathcal{C}) = \Lambda(y; a\mathcal{C}), \quad \forall a > 0. \quad (7)$$

5. Network iso-elasticity and weighted α -fairness

In this section, we prove the first result of this paper. We prove a network utility function satisfying this iso-elastic property is, up to an additive constant, a weighted α -fair utility function.

Theorem 5.1. For $U(y, \Lambda)$ the network utility function (1) the following are equivalent,

- i) $U(y, \Lambda)$ is network iso-elastic.
- ii) up to an additive constant, $U(y, \Lambda)$ is weighted α -fair, i.e., there exist $\alpha > 0$, $w \in \mathbb{R}_+^{\mathcal{R}}$ and constant $c \in \mathbb{R}$ such that

$$U(y, \Lambda) = W_{\alpha, w}(y, \Lambda) + c.$$

Proof. It is an immediate calculation that ii) implies i). We now prove conversely that i) implies ii). The key idea is to prove that the map $\phi_a : U(y, \Lambda) \mapsto U(y, a\Lambda)$ is linear.

For simplicity, we suppose $\sum_r y_r = \sum_r \tilde{y}_r = 1$; we, also, let $x_r = \frac{\Lambda_r}{y_r}$ and $\tilde{x}_r = \frac{\tilde{\Lambda}_r}{\tilde{y}_r}$ for $r \in \mathcal{R}$, and we let I and \tilde{I} be random variables with distribution y and \tilde{y} , respectively. One can see the value of the utility function

$$\mathbb{E}U_I(x_I) = \sum_{r \in \mathcal{R}} y_r U_r(x_r),$$

induces an ordering on the set of pairs (y, Λ) . Network iso-elasticity (4) states that, $\forall a > 0$, $\mathbb{E}U_I(x_I)$ and $\mathbb{E}U_I(ax_I)$ induce the same ordering on elements (I, x) . Thus (4) implies, for each $a > 0$, there exists an increasing function ϕ_a such that

$$\phi_a(\mathbb{E}U_I(x_I)) = \mathbb{E}U_I(ax_I), \quad \forall (I, x).$$

Now let us show that $\phi_a(\cdot)$ is linear. Let \hat{I} be the random variable such that

$$\hat{I} = \begin{cases} I & \text{with probability } p, \\ \tilde{I} & \text{with probability } (1 - p). \end{cases}$$

Let $\tau = \mathbb{E}U_I(x_I)$ and $\sigma = \mathbb{E}U_{\tilde{I}}(x_{\tilde{I}})$. Note $p\tau + (1 - p)\sigma = \mathbb{E}U_{\hat{I}}(x_{\hat{I}})$. Now observe

$$\begin{aligned} \phi_a(p\tau + (1 - p)\sigma) &= \mathbb{E}U_{\hat{I}}(ax_{\hat{I}}) = p\mathbb{E}U_I(ax_I) + (1 - p)\mathbb{E}U_{\tilde{I}}(ax_{\tilde{I}}) \\ &= p\phi_a(\tau) + (1 - p)\phi_a(\sigma). \end{aligned}$$

Thus $\phi_a(\cdot)$ is an increasing linear function and so, for $a > 0$,

$$\mathbb{E}U_I(ax_I) = b_a \mathbb{E}U_I(x_I) + d_a, \quad \forall (I, x), \quad (8)$$

for some $b_a > 0$ and $d_a \in \mathbb{R}$. Statement (8) applies to $I \equiv r$, for $r \in \mathcal{R}$. Thus

$$U_r(ax_r) = b_a U_r(x_r) + d_a, \quad \forall x_r \in \mathbb{R}_+.$$

We wish to remove the seemingly arbitrary functions b_a and d_a . Differentiating with respect to x_r gives,

$$U'_r(ax_r) = b_a U'_r(x_r), \quad \forall x_r \in \mathbb{R}_+. \quad (9)$$

As U'_r is a decreasing function it is differentiable at at least one point [6, Theorem 7.2.7]. Differentiating again at this point gives

$$U''_r(ax_r) = b_a U''_r(x_r), \quad \forall a > 0.$$

Thus, as a is arbitrary, $U''_r(x_r)$ exists for all $x_r > 0$. Dividing both terms leads to the equation

$$a \frac{U''_r(ax_r)}{U'_r(ax_r)} = \frac{U''_r(x_r)}{U'_r(x_r)}, \quad \text{for } x_r > 0, a > 0.$$

Taking $x_r = 1$ and defining $\alpha := -\frac{U_r''(1)}{U_r'(1)}$, we have the differential equation

$$a \frac{U_r''(a)}{U_r'(a)} = -\alpha_r, \quad \text{for } a > 0.$$

Dividing both sides by a and integrating with respect to a gives

$$\log U_r'(a) = -\alpha \log a + \log w_r \implies U_r'(a) = w_r a^{-\alpha_r}$$

Here $\log w_r$ is an appropriate additive constant. Integrating one more time gives and letting $a = x_r$ gives

$$U_r(x_r) = \begin{cases} w_r \frac{x_r^{1-\alpha_r}}{1-\alpha_r} + c_r & \text{if } \alpha_r \neq 1, \\ w_r \log x_r + c_r & \text{if } \alpha_r = 1. \end{cases} \quad (10)$$

It remains to show $\alpha_r = \alpha$ for all $r \in \mathcal{R}$. Observe that (4) includes the statement $U_r(x_r) \geq U_i(x_i)$ iff $U_r(ax_r) \geq U_i(ax_i)$, which can only hold if $\alpha_i = \alpha_r$. So $\alpha_r \equiv \alpha$ for some $\alpha > 0$ and thus for all $y \in \mathbb{R}_+^{\mathcal{R}}$ and $\Lambda \in \mathbb{R}_+^{\mathcal{R}}$

$$\begin{aligned} U(y, \Lambda) &= U\left(\frac{y}{\sum_i y_i}, \Lambda \sum_i y_i\right) \\ &= \begin{cases} \frac{1}{\sum_i y_i} \sum_{r \in \mathcal{R}} \frac{w_r y_r}{1-\alpha} \left(\frac{\Lambda_r}{y_r}\right)^{1-\alpha} & \text{if } \alpha \neq 1, \\ \frac{1}{\sum_i y_i} \sum_{r \in \mathcal{R}} w_r y_r \log\left(\frac{\Lambda_r}{y_r}\right) & \text{if } \alpha = 1, \end{cases} \end{aligned}$$

as required. \square

Remark 1. We note that an almost identical result and proof holds for the aggregate utility function

$$U(y, \Lambda) = \sum_{r \in \mathcal{R}} y_r U_r\left(\frac{\Lambda_r}{y_r}\right).$$

The only exception would be that the above result does not hold for the weighted proportionally fair case (where $\alpha = 1$).

6. Scaleability and weighted α -fairness in Networks with Local Traffic

In this section, we demonstrate that flow scalability is equivalent to optimizing a weighted α -fair utility function for two specific network topologies: a linear network and a network satisfying the local traffic condition.

A *linear network* consist of a routes r_0, \dots, r_K and links j_1, \dots, j_K for $K > 1$. Route r_0 uses all the links, $A_{jr_0} = 1 \ \forall j \in \mathcal{J}$, whilst for $k = 1, \dots, K$ each route r_k only uses link j_k , $A_{j_k r_k} = 1$ and $A_{jr_k} = 0$ otherwise. The *local traffic condition* insists that each link has a route that uses that link only, $\forall j \ \exists r$ such that $A_{jr} = 1$ and $A_{lr} = 0 \ \forall l \neq j$. We will also assume a network that satisfies

the local traffic condition has two or more links and is connected, that is for all links $j, l \in \mathcal{J}$ there exists a sequence of links and routes $j_0, r_0, j_1, r_1, \dots, r_{K-1}, j_K$ such that $j_0 = j$ and $j_K = l$, and $A_{j_k, r_k} = A_{j_{k+1}, r_k} = 1$ for $k = 0, \dots, K-1$.

We now prove that flow scalability is equivalent to weighted α -fairness for linear networks.

Theorem 6.1. *Assuming, for $r \in \mathcal{R}$, $U'_r(x)$ has range $(0, \infty)$, a linear network is flow scalable iff it maximizes a weighted α -fair utility function.*

Proof. It is immediate that the weighted α -fair utility is flow scalable. Let us show the converse.

The bandwidth a linear network allocates is the solution to the optimization

$$\max \quad \sum_{r \in \mathcal{R}} y_r U'_r \left(\frac{\Lambda_r}{y_r} \right) \quad (11)$$

$$\text{subject to} \quad \Lambda_{r_0} + \Lambda_{r_k} \leq C_{r_k}, \quad k = 1, \dots, K. \quad (12)$$

Here, we let C_{r_k} denote the capacity of the link on the intersection of route r_k and route r_0 . Assuming $y_r > 0 \forall r \in \mathcal{R}$, the solution to this optimization satisfies

$$C_{r_k} = \Lambda_{r_0} + \Lambda_{r_k}, \quad k = 1, \dots, K \quad (13)$$

$$\text{and} \quad U'_{r_0} \left(\frac{\Lambda_{r_0}}{y_{r_0}} \right) = \sum_{k=1}^K U'_{r_k} \left(\frac{\Lambda_{r_k}}{y_{r_k}} \right). \quad (14)$$

If the solution is flow scalable then Λ must also satisfy

$$U'_{r_0} \left(\frac{\Lambda_{r_0}}{ay_{r_0}} \right) = \sum_{k=1}^K U'_{r_k} \left(\frac{\Lambda_{r_k}}{ay_{r_k}} \right), \quad \forall a > 0. \quad (15)$$

We observe that the solutions for a linear network allow a suitably large array of values for $U'_r \left(\frac{\Lambda_r}{y_r} \right)$ and Λ_r . In particular, for any $p_1, \dots, p_K \in (0, \infty)$ and $\Lambda_{r_0}, \dots, \Lambda_{r_K} \in (0, \infty)$, we can choose a $y_r > 0$ such that $p_k = U'_{r_k} \left(\frac{\Lambda_{r_k}}{y_{r_k}} \right)$, $k = 1, \dots, K$. Letting $p_0 = \sum_{k=1}^K p_k$, we can also choose y_{r_0} such that $p_0 = U'_{r_0} \left(\frac{\Lambda_{r_0}}{y_{r_0}} \right)$. Thus, this choice of Λ gives the unique solution to (11) for the y given.

The remainder of the proof is similar to that of Theorem 5.1. Once again, the key idea is to prove the map $\phi_a^r : U'_r \left(\frac{\Lambda_r}{y_r} \right) \mapsto U'_r \left(\frac{\Lambda_r}{ay_r} \right)$ is linear. As $U'_r(x)$ must be continuous and decreasing $\phi_a^r(\cdot)$ is an increasing continuous function. Also, note that $\phi_a^r(0) = 0$ because $U'_r(\infty) = 0$.

With $p_r = U'_r \left(\frac{\Lambda_r}{y_r} \right)$ as above, (14) and (15) state that

$$\phi_a^{r_0}(p_0) = \sum_{k=1}^K \phi_a^{r_k}(p_k) \quad \text{for} \quad p_0 = \sum_{k=1}^K p_k. \quad (16)$$

Taking $p_0 = p$, $p_k = p$ and $p_{k'} = 0$ for $k' \neq k$, we see that

$$\phi_a^{r_0}(p) = \phi_a^{r_k}(p), \quad \forall p \geq 0. \quad (17)$$

Taking $p_0 = p$, $p_1 = \hat{p}$ and $p_2 = p - \hat{p}$ for $p > \hat{p} \geq 0$, (16) and (17) together imply

$$\frac{\phi_a^r(p) - \phi_a^r(\hat{p})}{p - \hat{p}} = \frac{\phi_a^{r_0}(p - \hat{p})}{p - \hat{p}}. \quad (18)$$

As $\phi_a^{r_0}$ is increasing it is differentiable at some point (see [6, Theorem 7.2.7]), then by (18) it is differentiable at all points and moreover its derivative is constant. In other words, ϕ_a^r must be linear,

$$\phi_a^r(p_r) = b_a p_r$$

for some function b_a . Thus

$$U_r'(ax_r) = b_a U_r'(x_r), \quad \forall x_r \in \mathbb{R}_+.$$

By same argument used to derive (10) from (9) in Theorem 5.1, we have that

$$U_r(x_r) = \begin{cases} w_r \frac{x_r^{1-\alpha_r}}{1-\alpha_r} + c_r & \text{if } \alpha_r \neq 1, \\ w_r \log x_r + c_r & \text{if } \alpha_r = 1. \end{cases}$$

$\alpha_r > 0$, $w_r > 0$ and $c_r \in \mathcal{R}$.

Finally note, if $\alpha_{r_k} \neq \alpha_{r_0}$ for some r_k then the equalities (15) and (14) cannot hold, for all $a > 0$. Thus, we now see the optimization (11-12) must have been a weighted α -fair optimization. \square

From this argument for linear networks, we extend our result to any network satisfying the local traffic condition.

Corollary 6.1. *Assuming, for $r \in \mathcal{R}$, $U_r'(x)$ has range $(0, \infty)$, a network satisfying is the local traffic condition is flow scalable iff it maximizes a weighted α -fair utility function.*

Proof. It is immediate that the weighted α -fair utility is flow scalable. To show the converse, we show a local traffic network can be reduced to a linear network. Take any route $r_0 \in \mathcal{R}$ which uses two or more links. Let them be j_1, \dots, j_K and let r_1, \dots, r_K be local traffic routes for each respective link. Set $y_r = 0$ for all $r \notin \{r_0, \dots, r_K\}$. The resulting network is a linear network, and thus, by Theorem 6.1, the utility function associated with each route of this subnetwork is of the form

$$U_r(x_r) = \begin{cases} w_r \frac{x_r^{1-\alpha_r}}{1-\alpha_r} + c_r & \text{if } \alpha_r \neq 1, \\ w_r \log x_r + c_r & \text{if } \alpha_r = 1. \end{cases}$$

$\alpha_r > 0$, $w_r > 0$ and $c_r \in \mathcal{R}$. Here α_r is same for all routes in this linear subnetwork, i.e. $\alpha_{r_k} = \alpha_{r_0}$ for all $k = 1, \dots, K$.

We show, that since our network is connected, α_r does not depend on the route chosen. We know $\alpha_r = \alpha_{r(j)}$ for each local traffic route $r(j)$ intersecting

route r . Take two routes r and \tilde{r} . Since our network is connect, for j a link on route r and l a link on route \tilde{r} , there are a sequence $j_0, r_0, j_1, r_1, \dots, r_{K-1}, j_K$ connecting $j_0 = j$ and $j_K = l$. For each $k = 0, \dots, k-1$, the subnetwork formed by r_k, j_k and j_{k+1} and their respective local traffic routes forms a linear network. Thus α_r is constant across all routes and associated local traffic routes in the sequences $j_0, r_0, j_1, r_1, \dots, r_{K-1}, j_K$. This, thus, implies $\alpha_r = \alpha_{\tilde{r}}$. \square

By simply adding a local traffic route to each link, any network topology can be made to satisfy the local traffic condition. Thus given the generality of this set of networks, it is natural to conjecture that flow scalability is equivalent to weighted α -fairness for any network topology.

Conjecture 6.1. *A network is flow scalable iff it maximizes a weighted α -fair utility function.*

7. Access scalability and weighted α -fairness

The network utility maximization problem is

$$\max U(y, \Lambda) \quad \text{over} \quad \Lambda \in \mathcal{C}. \quad (19)$$

Given the flows on each route y , this finds an optimal way to allocate bandwidth Λ . Given the bandwidth to be allocated is Λ , what is the optimal number of flows permissible on each route? This leads to the maximization

$$\max U(y, \Lambda) \quad \text{over} \quad y \in \mathcal{Y}. \quad (20)$$

In this second optimization, we optimize the number of flows accessing routes inorder to guarantee the maximum average utility per flow.

Flow scalability (6-7) says, if we multiply the capacity of the network then the rate allocated to each route will multiply by the same amount. Similarly, we will say a utility function is access scalable, if multiplying the bandwidth available to each route means that the number flows accepted on each route will multiply by the same amount. More formally,

Definition 3. *We say $y(\Lambda)$, the solution to optimization problem (20), is access scalable if, for all $a > 0$,*

$$y(a\Lambda) = ay(\Lambda).$$

In the following theorem, we observe that the only utility function that is both flow scalable and access scalable is a weighted α -fair utility function.

Theorem 7.1. *A network utility function $U(y, \Lambda)$ is both flow scalable for all \mathcal{C} and access scalable for all \mathcal{Y} iff it is weighted α -fair.*

Proof. First, it is an immediate calculation that any weighted α -fair utility function is user scalable and access scalability for any set \mathcal{C} and \mathcal{Y} .

Conversely, suppose user scalability and access scalability hold for each \mathcal{C} and \mathcal{Y} . Taking $\mathcal{C} = \{\Lambda, \tilde{\Lambda}\}$ and $\mathcal{Y} = \{y, \tilde{y}\}$, by user scalability and access scalability, respectively, we have that

$$U(y, \Lambda) \leq U(y, \tilde{\Lambda}) \quad \text{iff} \quad U(ay, \Lambda) \leq U(ay, \tilde{\Lambda}), \quad (21)$$

$$U(y, \Lambda) \leq U(\tilde{y}, \Lambda) \quad \text{iff} \quad U(ay, \Lambda) \leq U(a\tilde{y}, \Lambda). \quad (22)$$

By continuity, there exists an \tilde{a} such that

$$U(\tilde{y}, \tilde{\Lambda}) = U(\tilde{a}\tilde{y}, \Lambda). \quad (23)$$

Also, by (21),

$$U(\tilde{y}, \tilde{\Lambda}) = U(\tilde{a}\tilde{y}, \Lambda) \quad \text{iff} \quad U(a\tilde{a}\tilde{y}, \Lambda) = U(ay, \tilde{\Lambda}). \quad (24)$$

Now, (22–24) imply the following set of equivalences.

$$\begin{aligned} U(y, \Lambda) &\leq U(\tilde{y}, \tilde{\Lambda}) = U(\tilde{a}\tilde{y}, \Lambda) && \text{[by (23)]} \\ \text{iff} \quad U(ay, \Lambda) &\leq U(a\tilde{a}\tilde{y}, \Lambda) && \text{[by (22)]} \\ \text{iff} \quad U(ay, \Lambda) &\leq U(a\tilde{y}, \tilde{\Lambda}). && \text{[by (24)]} \end{aligned}$$

Thus, $U(y, \Lambda)$ is network iso-elastic and so, by Theorem 5.1, $U(y, \Lambda)$ is a weighted α -fair utility function. \square

References

- [1] ARROW, J. K. (1965). *Aspects of the theory of risk-bearing: Yrj Jahnsson lectures*. Yrj Jahnssonin Sti, New York.
- [2] BERTSEKAS, D. AND GALLAGER, R. (1987). *Data Networks*. Prentice Hall.
- [3] BONALD, T. AND MASSOULIÉ, L. (2001). Impact of fairness on internet performance. *Proc. of ACM Sigmetrics 29*, 82–91.
- [4] CHERRY, S. (2004). Telecom: Edholm’s law of bandwidth. *IEEE Spectr.* **41**, 7, 58–60.
- [5] COFFMAN, K. G. AND ODLYZKO, A. M. (2002). Internet growth: Is there a moores law for data traffic? In *Handbook of Massive Data Sets*. Kluwer, 47–93.
- [6] DUDLEY, R. M. (2002). *Real Analysis and Probability*, 2nd ed. Cambridge University Press.
- [7] EGOROVA, R., BORST, S., AND ZWART, B. (2007). Bandwidth-sharing networks in overload. *Performance Evaluation* **64**, 9-12, 978 – 993. Performance 2007, 26th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation.
- [8] GANESH, A., LAEVEN, K., AND STEINBERG, R. (2007). Congestion pricing and noncooperative games in communication networks. *Oper. Res.* **55**, 3, 430–438.

- [9] GROMOLL, H. C. AND WILLIAMS, R. J. Fluid limits for networks with bandwidth sharing and general document size distributions. *Ann. Appl. Probab.* **19**, 1.
- [10] KANG, W. N., KELLY, F. P., LEE, N. H., AND WILLIAMS, R. J. (2009). State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. *Ann. Appl. Probab.* **19**, 1719–1780.
- [11] KELLY, F. P. (1997). Charging and rate control for elastic traffic. *European Transactions on Telecommunications* **8**, 33–37.
- [12] KELLY, F. P., MAULLOO, A. K., AND TAN, D. K. H. (1998). Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* **49**, 237–252.
- [13] LA, R. AND MO, J. (2009). Interaction of service providers in task delegation under simple payment rules. *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, 8594–8599.
- [14] LAN, T., KAO, D., CHIANG, M., AND SABHARWAL, A. (2010). An axiomatic theory of fairness in network resource allocation. In *INFOCOM, 2010 Proceedings IEEE*. 1–9.
- [15] LOW, S. H. (2000). A duality theory of tcp flow controls. *Proceedings of at ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*.
- [16] MASSOULIÉ, L. AND ROBERTS, J. (1999). Bandwidth sharing: Objectives and algorithms. *IEEE Infocom 1999* **10**, 3, 320–328.
- [17] MO, J. AND WALRAND, J. (2000). Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Networking* **8**, 556–567.
- [18] ODLYZKO, A. M. (2008). Threats to the internet: Too much or too little growth? *Internet Evolution*. February.
- [19] PRATT, J. (1964). Risk aversion in the small and in the large. *Econometrica* **32**, 122–136.
- [20] UCHIDA, M. AND KUROSE, J. (2009). An information-theoretic characterization of weighted alpha-proportional fairness. In *INFOCOM 2009, IEEE*. 1053–1061.
- [21] YE, H.-Q. AND YAO, D. D. (2010). Utility-maximizing resource control: Diffusion limit and asymptotic optimality for a two-bottleneck model. *Operations Research* **58**, 3, 613–623.